Parallel implementation of FETI-2LM for large problems with many RHS in CEM

François-Xavier Roux
Roberto Molina (LJLL, UPMC)

ONERA
THE FRENCH AEROSPACE LAB

return on innovation

Summary

- Antenna array problem
- FETI-2LM
- FETI-2LM for Maxwell with Nédélec finite elements
- Acceleration of iterations via reuse of search directions
- Parallel local direct solver
- Block strategy
- Conclusion

Antenna Arrays

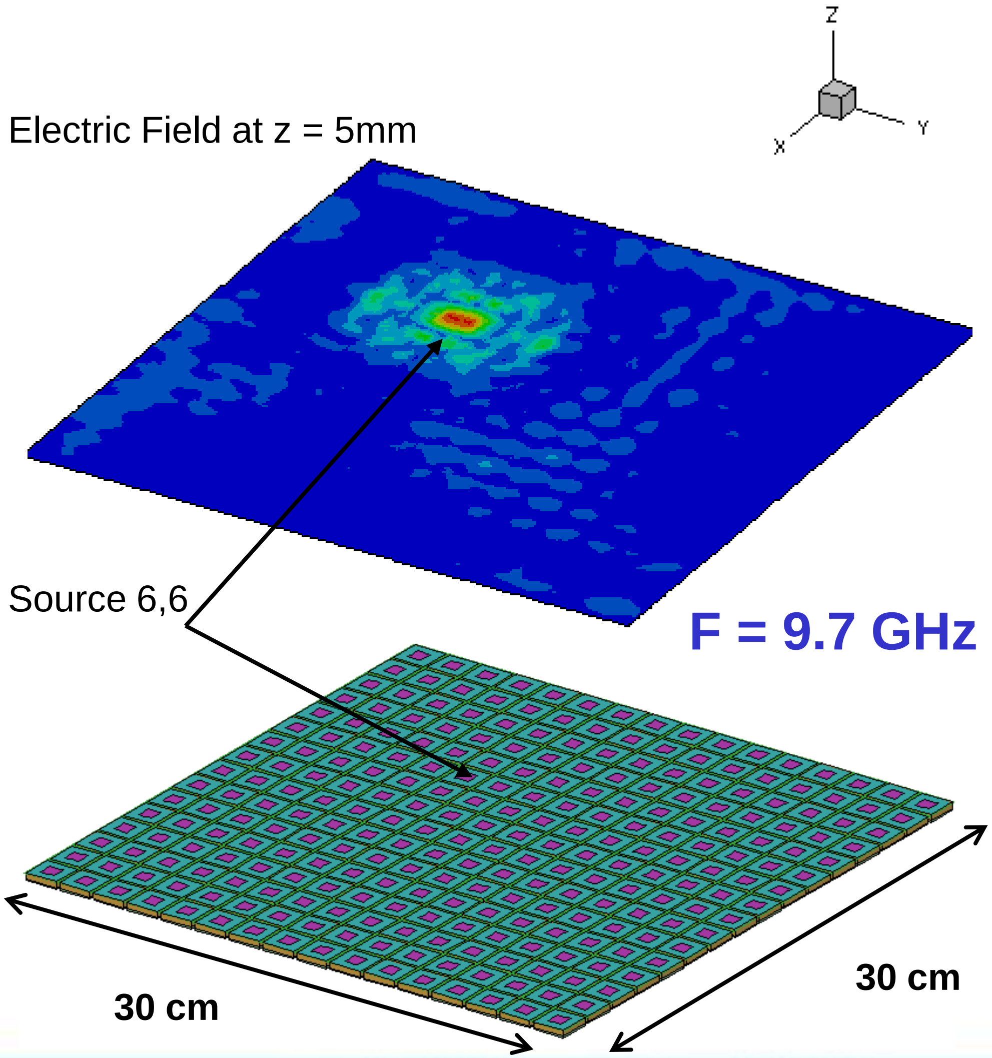

- 2D device, attached to the fuselage of the plane
- Complex treatment to orient the radar beam
- Larger arrays for higher intensity

- Fine local mesh to take into account the complex structure of each element

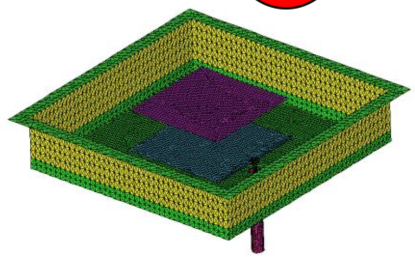

Maxwell equations

Find the diffracted electric field $\vec{E} \in H(curl;\Omega)$ such that

Isotropic & anisotropic Dielectric materials

$$
\begin{cases}
\nabla \times \left(\dfrac{1}{\vec{\mu}_r} \nabla \times \vec{E}\right) - k_0^2 \vec{\varepsilon}_r \vec{E} = k_0^2 (\varepsilon_{r,i} - \mu_{r,i}^{-1}) \vec{E}_{incident} & in \quad \Omega \subset R^3 \\
Z_0 R_e (\vec{n} \times \vec{H}) + \vec{n} \times (\vec{n} \times \vec{E}) = 0 & on \quad \Gamma_r \\
\vec{n} \times \vec{E} = -\vec{n} \times \vec{E}_{incident} & on \quad \Gamma_e \, (PEC) \\
\vec{n} \times (\nabla \times \vec{E}) + jk_0 \vec{n} \times (\vec{n} \times \vec{E}) = 0 & on \quad \Gamma_{ext}
\end{cases}
$$

ABC

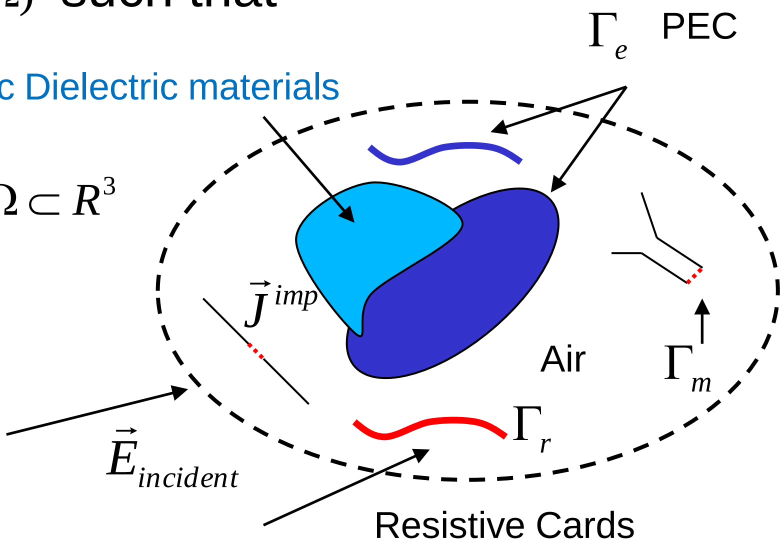

Multiple sources:

- plane waves
- modes

FEM: Nédélec tetrahedral elements

- Degree 1 (6 dof)
- Degree 2 (20 dof)

Domain decomposition method with Robin interface conditions

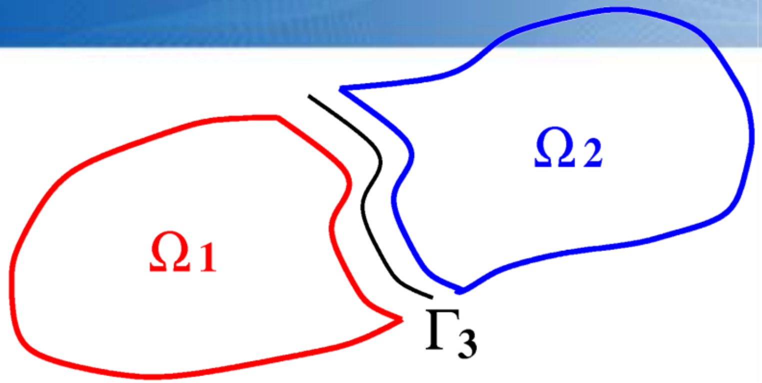

- Local Problems with Robin boundary conditions

$$\begin{cases} -\Delta u_1 = f \quad \text{in } \Omega_1 \\ \dfrac{\partial u_1}{\partial n_1} + \alpha_1 u_1 = \lambda_1 \text{ on } \Gamma_3 \end{cases}$$

$$\begin{cases} -\Delta u_2 = f_2 \quad \text{in } \Omega_2 \\ \dfrac{\partial u_2}{\partial n_2} + \alpha_2 u_2 = \lambda_2 \text{ on } \Gamma_3 \end{cases}$$

- Matching conditions on Γ_3

$$\begin{cases} u_1 - u_2 = 0 \\ \dfrac{\partial u_1}{\partial n_1} + \dfrac{\partial u_2}{\partial n_2} = 0 \end{cases} \Leftrightarrow$$

$$\begin{cases} u_1 - u_2 = 0 \\ \lambda_1 - \alpha_1 u_1 + \lambda_2 - \alpha_2 u_2 = 0 \end{cases} \Leftrightarrow$$

$$\begin{cases} \lambda_1 - \alpha_1 u_1 + \lambda_2 - \alpha_2 u_1 = 0 \\ \lambda_1 - \alpha_1 u_2 + \lambda_2 - \alpha_2 u_2 = 0 \end{cases}$$

Application to discrete problems : FETI-2LM method

- Global system of equations

$$\begin{pmatrix} K_{11} & 0 & K_{13} \\ 0 & K_{22} & K_{23} \\ K_{31} & K_{32} & K_{33} \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} = \begin{pmatrix} b_1 \\ b_2 \\ b_3 \end{pmatrix}$$

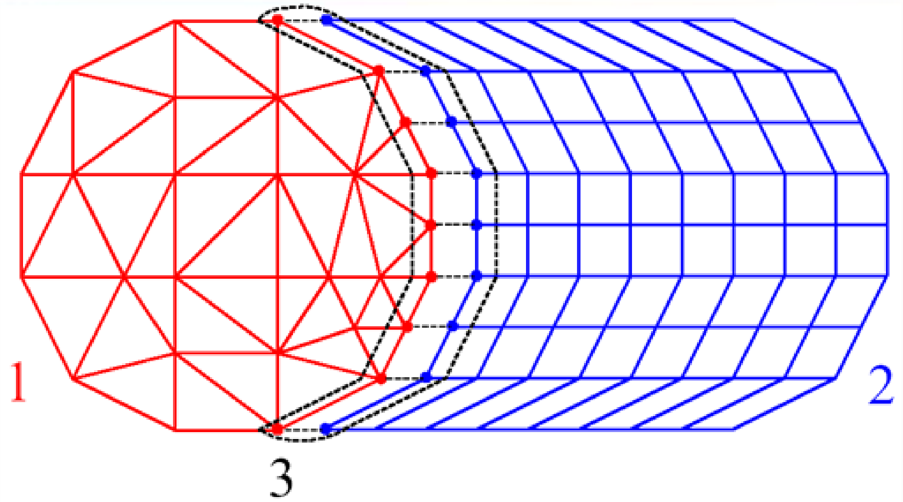

- Interface conditions

$$\begin{cases} x_3^{(1)} = x_3^{(2)} \\ K_{31}x_1 + K_{32}x_2 + K_{33}^{(1)}x_3^{(1)} + K_{33}^{(2)}x_3^{(2)} = b_3 \\ \Leftrightarrow \lambda_1 + \lambda_2 - k_1 x_3^{(1)} - k_2 x_3^{(2)} = 0 \end{cases}$$

- Local system of equations

$$\begin{pmatrix} K_{11} & K_{13} \\ K_{31} & K_{33}^{(1)} + k_1 \end{pmatrix} \begin{pmatrix} x_1 \\ x_3^{(1)} \end{pmatrix} = \begin{pmatrix} b_1 \\ b_3^{(1)} + \lambda_1 \end{pmatrix}$$

$$\begin{pmatrix} K_{22} & K_{23} \\ K_{32} & K_{33}^{(2)} + k_2 \end{pmatrix} \begin{pmatrix} x_2 \\ x_3^{(2)} \end{pmatrix} = \begin{pmatrix} b_2 \\ b_3^{(2)} + \lambda_2 \end{pmatrix}$$

$$\begin{cases} \lambda_1 + \lambda_2 - (k_1 + k_2)\, x_3^{(2)} = 0 \\ \lambda_1 + \lambda_2 - (k_2 + k_1)\, x_3^{(1)} = 0 \end{cases}$$

Condensed interface problem

- Local condensation

$$\left(k_1 + K_{33}^{(1)} - K_{31} K_{11}^{-1} K_{13}\right) x_3^{(1)} = \lambda_1 + b_3^{(1)} - K_{31} K_{11}^{-1} b_1$$

$$\left(k_2 + K_{33}^{(2)} - K_{32} K_{22}^{-1} K_{23}\right) x_3^{(2)} = \lambda_2 + b_3^{(2)} - K_{32} K_{22}^{-1} b_2$$

- Matrix of interface problem

$$\begin{pmatrix} I & I - (k_1 + k_2)\left(k_2 + K_{33}^{(2)} - K_{32} K_{22}^{-1} K_{23}\right)^{-1} \\ I - (k_2 + k_1)\left(k_1 + K_{33}^{(1)} - K_{31} K_{11}^{-1} K_{13}\right)^{-1} & I \end{pmatrix}$$

Optimal Robin interface conditions

- Optimal interface conditions

$$k_1 = K_{33}^{(2)} - K_{32} K_{22}^{-1} K_{23}$$
$$k_2 = K_{33}^{(1)} - K_{31} K_{11}^{-1} K_{13}$$

- Optimal conditions = static condensation on interface of remaining structure
- Interpretation via local static condensation in global system of equations

$$\begin{pmatrix} K_{11} & K_{13} \\ K_{31} & K_{33}^{(1)} + K_{33}^{(2)} - K_{32} K_{22}^{-1} K_{23} \end{pmatrix} \begin{pmatrix} x_1 \\ x_3 \end{pmatrix} = \begin{pmatrix} b_1 \\ b_3^{(1)} + b_3^{(2)} - K_{32} K_{22}^{-1} b_2 \end{pmatrix}$$

- Dirichlet to Neumann mapping for domain exterior

FETI-2LM applied to Maxwell

$\Omega = \Omega_1 \cup \Omega_2 .... \cup \Omega_N$ Domain partition

$$\begin{cases} \nabla \times (\frac{1}{\vec{\mu}_r} \nabla \times \vec{E}_i) - k_0^2 \vec{\varepsilon}_r \vec{E}_i = k_0^2 (\varepsilon_{r,i} - \mu_{r,i}^{-1}) \vec{E}_{incident} & in \quad \Omega_i \subset R^3 \\ \vec{n}_i \times (\frac{1}{\vec{\mu}_{r,i}} \nabla \times \vec{E}_i) + jk_0 \vec{n}_i \times (\vec{n}_i \times \vec{E}_i) = \vec{\Lambda}^i_j & on \quad \Gamma_i \quad (Robin) \\ \vec{n} \times (\nabla \times \vec{E}_i) + jk_0 \vec{n} \times (\vec{n} \times \vec{E}_i) = 0 & on \quad \Gamma_{ext} = \partial\Omega_i \setminus \Gamma_i \end{cases}$$

Robin

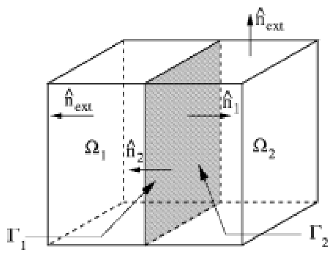

Additional variables on the interface (2 LM method)

Lagrange multipliers with approximate transparent Robin condition (approximate outer Dirichlet-Neumann)

$$\vec{n}_i \times (\mu_{r,i}^{-1} . \nabla \times \vec{E}^i_j) + jk_0 \vec{n}_i \times (\vec{n}_i \times \vec{E}^i_j) = \vec{\Lambda}^i_j$$

$$\vec{n}_j \times (\mu_{r,j}^{-1} . \nabla \times \vec{E}^j_i) + jk_0 \vec{n}_j \times (\vec{n}_j \times \vec{E}^j_i) = \vec{\Lambda}^j_i$$

Electric and Magnetic field Continuity

$$\vec{n}_i \times (\vec{n}_i \times \vec{E}^i_j) = \vec{n}_j \times (\vec{n}_j \times \vec{E}^j_i) \qquad (1)$$

$$\vec{n}_i \times (\mu_{r,i}^{-1} . \nabla \times \vec{E}^i_j) = -\vec{n}_j \times (\mu_{r,j}^{-1} . \nabla \times \vec{E}^j_i) \qquad (2)$$

$$\Lambda^i_j + \Lambda^j_i - 2jk_0 \vec{n}_i \times (\vec{n}_i \times \vec{E}^i_j) = 0$$

$$\Lambda^i_j + \Lambda^j_i - 2jk_0 \vec{n}_j \times (\vec{n}_j \times \vec{E}^j_i) = 0$$

Krylov subspace method and iterations

$$\begin{cases} \Lambda^i_j + \Lambda^j_i - 2jk_0 \vec{n}_i \times \vec{n}_i \times E^i_j = 0 \\ \Lambda^i_j + \Lambda^j_i - 2jk_0 \vec{n}_j \times \vec{n}_j \times E^j_i = 0 \end{cases} \quad on \quad \Gamma^{ij}$$

$$\lambda^i_j + \lambda^j_i - (M^i_j + M^j_i)E^i_j = 0 \qquad i = 1,2,..,N_s \quad and \quad j \in neighbor(i)$$

$$M^i_j = jk_0 \int_{\Gamma_{ij}} (\vec{n}_i \times \vec{W}_i).(\vec{n}_i \times \vec{W}_i)dS$$

Computation of $F\lambda - d$

1. Solution of local problem with Robin conditions defined by λ
2. Exchange values of E and λ on interfaces
3. On each interface Γ_i^j, computation of :

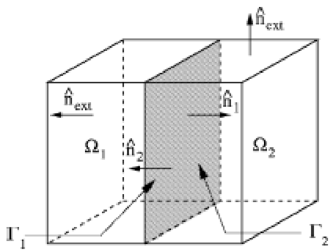

$$g^j_i = \lambda^i_j + \lambda^j_i - (M^i_j + M^j_i)E^i_j$$

ORTHODIR iterations until || K x - b || < stopping criterion

ONERA

Tests on prototype antenna array

- 44x44 array
- 1936 subdomains, 400M unknowns
- Comparison between numerical results and experimental results OK

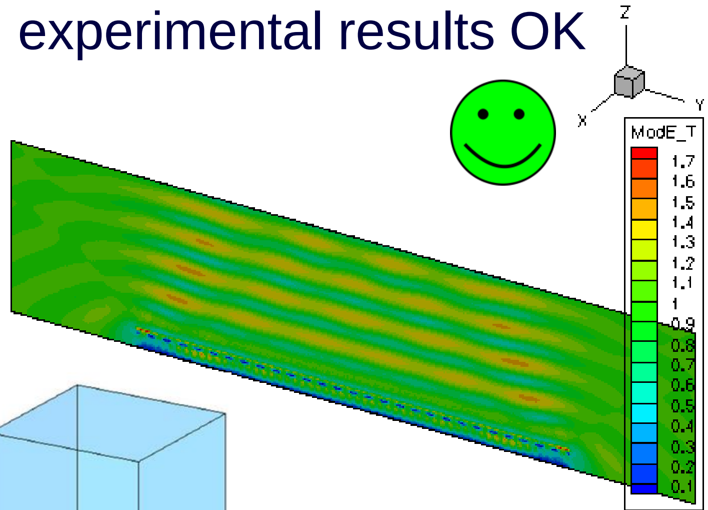

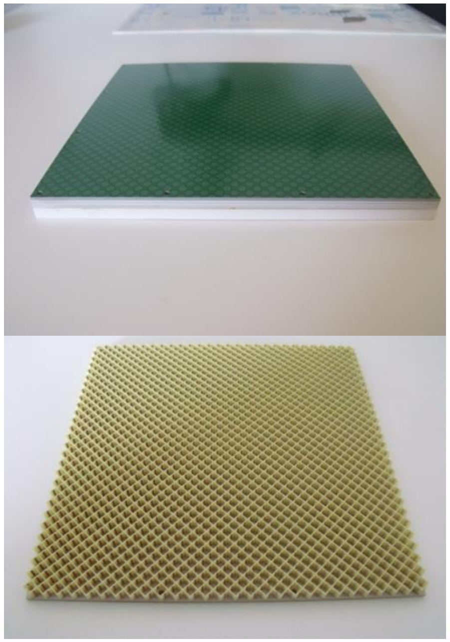

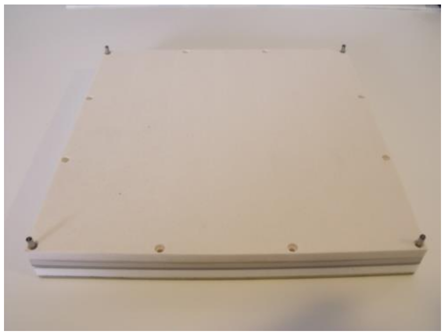

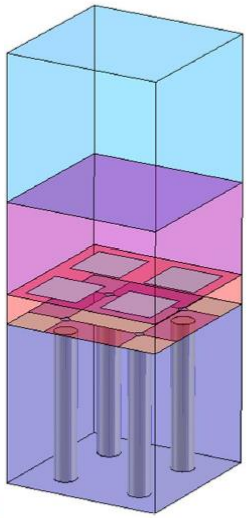

Weak scalability

- Timings with one subdomain per core
- Increasing size of the array

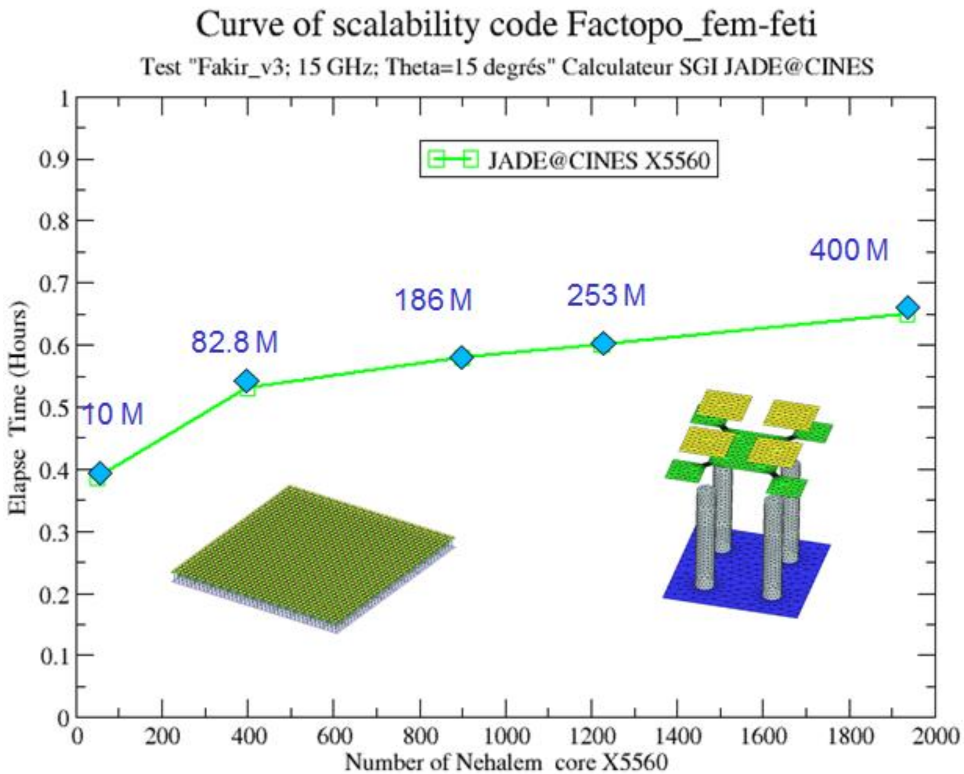

ORTHODIR algorithm

- ORTHODIR : build a *F*F*-orthogonal basis of Krylov space

$$\begin{cases} g_0 = F\lambda_0 - d \\ v_0 = g_0 \\ Fv_0 \\ d_0 = \left\| Fv_0 \right\| \end{cases}$$

$$\begin{cases} \lambda_p = \lambda_{p-1} + \rho_{p-1} v_{p-1} \\ g_p = g_{p-1} + \rho_{p-1} F v_{p-1} \\ (Fv_{p-1})^* g_p = 0 \Leftrightarrow d_{p-1}\rho_{p-1} = -(Fv_{p-1})^* g_{p-1} \\ \qquad\qquad\qquad\qquad\qquad\qquad - (Fv_{p-1})^* g_0 \end{cases}$$

$$\begin{cases} v_p = Fv_{p-1} + \sum_0^{p-1} \gamma_{ip} v_i \\ Fv_p = FFv_{p-1} + \sum_0^{p-1} \gamma_{ip} Fv_i \\ \gamma_{ip} = -(Fv_i)^* FFv_{p-1} \\ d_p = \left\| Fv_p \right\| \end{cases}$$

- ORTHODIR : block formulation

$$V^p = \left[v_0 v_1 \cdots v_{p-1} \right], \; \left(FV^p \right)^* \left(FV^p \right) = D^p, \; D^p \rho^p = -\left(FV^p \right)^* g_0$$

$$\begin{cases} \lambda_p = \lambda_0 + V^p \rho^p \\ g_p = g_0 + FV^p \rho^p \end{cases} \qquad \rho^p = -\left(FV^p \right)^* g_0$$

Restarted ORTHODIR with multiple RHS

- V^p and $F\,V^p$ are given , $(F\,V^p)^*\,(F\,V^p) = D_p$
- Optimal starting λ_0^{opt}

$$\begin{cases} \lambda_0^{opt} = \lambda_0 + V^p \rho^p \\ g_0^{opt} = g_0 + FV^p \rho^p \end{cases} \qquad D^p \rho^p = -\left(FV^p\right)^* g_0 \Leftrightarrow -\left(FV^p\right)^* g_0^{opt} = 0$$

- Start new iterations with new search directions *F*F*-orthogonal to V^p
- *F*F*-projected ORTHODIR
- In practice same as if restarting ORTHODIR at iteration *p*
- Accumulation of search directions with successive RHS

Dependancy upon number of stored directions

- 17x17 array
- 289 subdomains, 50 Million unknowns
- 31 RHS, incident waves with various angles

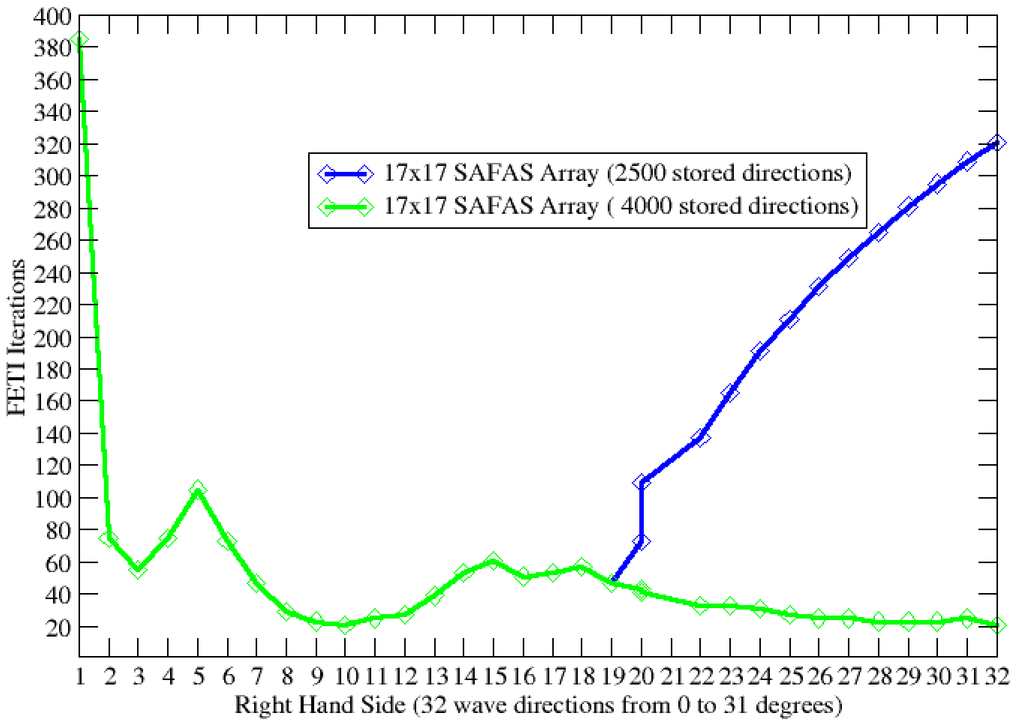

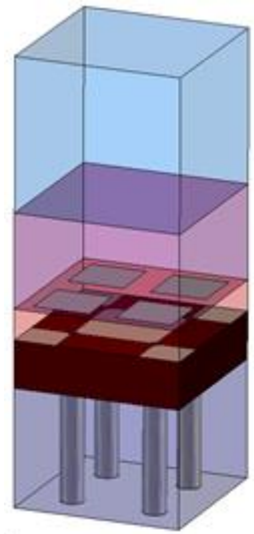

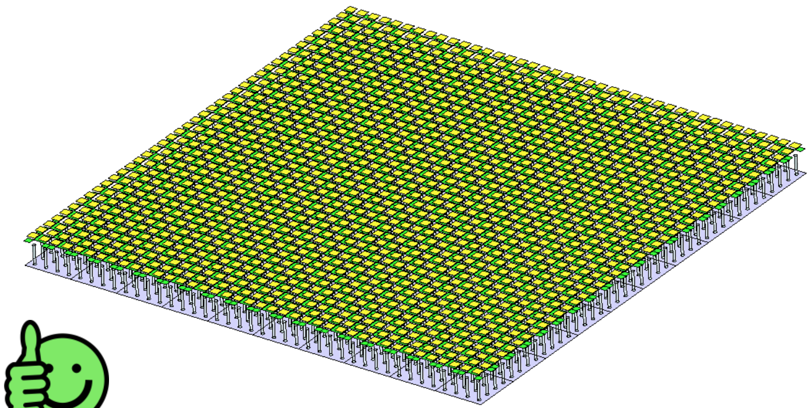

Computer architecture

- Networked compute nodes
- Each node is a hierarchical memory SMP with possibly parallel co-processor
- Space and time locality of data required for performance
- Present trend: increasing number of cores on each node

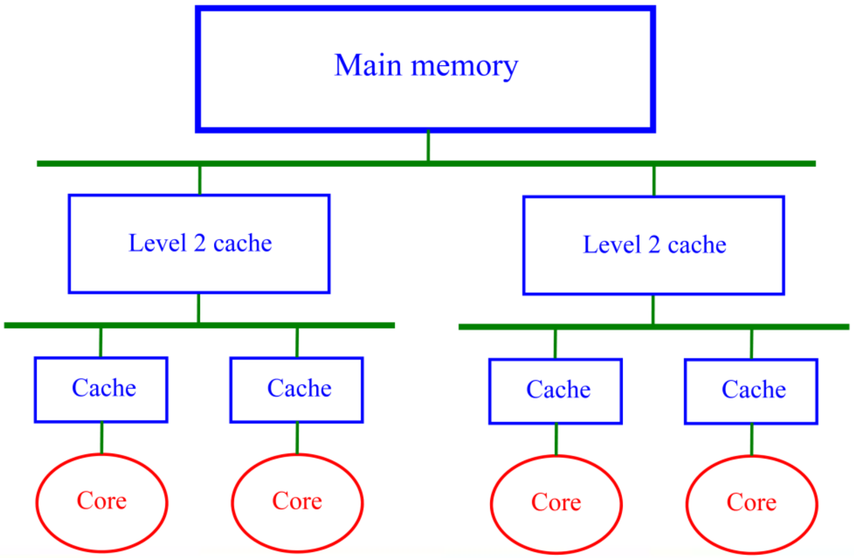

Local direct solver on SMP compute node

- Sparse direct solver based on nested bisection
- P-threads multi-threading management
- Splitting of blocks in small sub-blocks

	Dissection		Pardiso	
# cores	CPU time	elapsed time	CPU	esapsed time
1	74.84	72.824	85.04	82.941
2	74.81	38.162	87.79	43.627
4	77.32	20.454	92.66	23.141
6	79.96	15.200	104.38	17.391
8	83.56	12.008	118.25	14.786
12	94.08 (×1.26)	9.873 (/7.38)	165.99 (×1.95)	13.993 (/5.93)

`elstct1`, $N = 206,763$, nonsingular

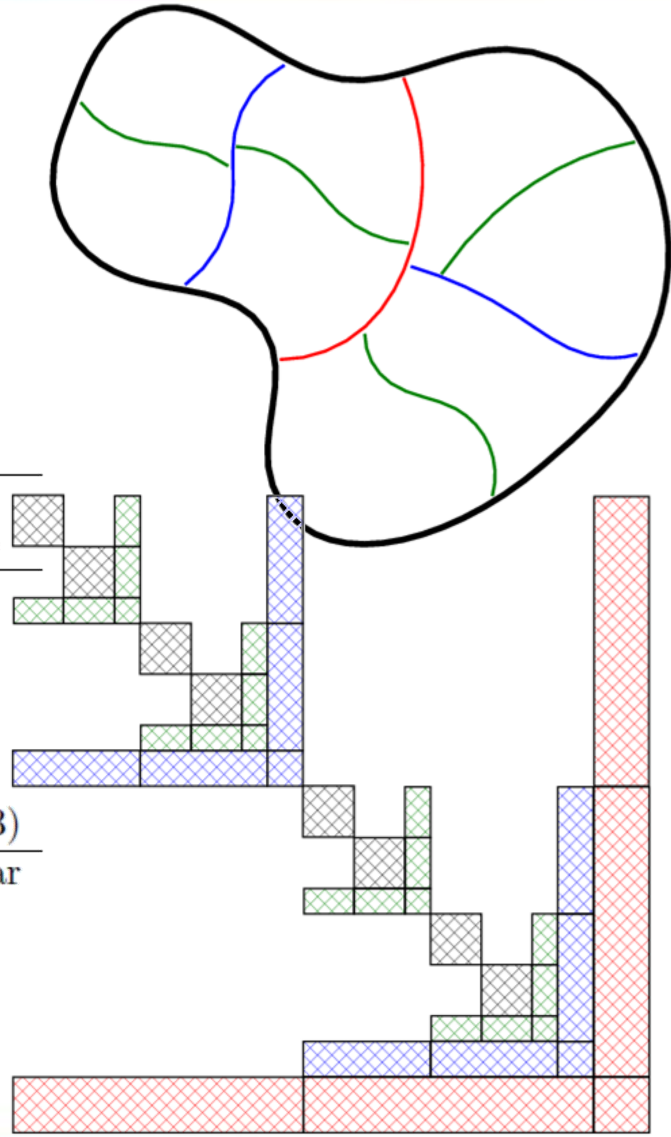

Efficiency of multiple forward-backward substitution

- Actual performance limited by global memory access
- Parallelization of forward-backward substitution for a single RHS gives very limited performance (speed-up < 2)
- With multiple RHS, higher arithmetic complexity with same memory access requirement (better data locality)

	1RHS @ 1core	12 RHS @ 12core	efficiency
Dissection	0.6194 sec.	0.5135 sec	120.6%
Pardiso	0.7054 sec.	1.2642 sec	55.8%

- With more than one subdomain per node, memory bandwidth available for each MPI process is even lower
- Performance of each single RHS forward-backward substitution is even poorer

Limited parallel efficiency of restarted ORTHODIR

- One product by *F* per iteration
- Single RHS local forward-backward substitution

- For numerical stability, use modified Gram-Schmidt procedure for orthogonalization
- dot product (BLAS1) + global reduction via MPI one by one

- Inefficient for local multi-threading on multi-core node
- Large communication overhead

- Use simultaneous solution
- Keep good properties of restarted ORTHODIR

} => block ORTHODIR

Block ORTHODIR algorithm

- Block ORTHODIR initialization

$$\begin{cases} g_0^k = F\lambda_0^k - d^k \\ v_0^k = g_0^k \end{cases}, \quad k = 1, n_{block}$$

$$\begin{cases} \Lambda_0 = \left[\lambda_0^1 \lambda_0^2 \cdots \lambda_0^{n_{block}}\right] \\ G_0 = \left[g_0^1 g_0^2 \cdots g_0^{n_{block}}\right] \\ V_0 = \left[v_0^1 v_0^2 \cdots v_0^{n_{block}}\right] \\ D_0 = (FV_0)^*(FV_0) = L_0 L_0^* \end{cases}$$

- Block ORTHODIR iteration

$$\begin{cases} \Lambda_p = \Lambda_{p-1} + V_{p-1} \mathrm{P}_{p-1} \\ G_p = G_{p-1} + FV_{p-1} \mathrm{P}_{p-1} \\ FV_{p-1})^* G_p = 0 \Leftrightarrow D_{p-1} \mathrm{P}_{p-1} = -(FV_{p-1})^* G_{p-1} \end{cases}$$

$$\begin{cases} V_p = FV_{p-1} + \sum_0^{p-1} V_i \, \Gamma_{ip} \\ FV_p = FFV_{p-1} + \sum_0^{p-1} FV_i \, \Gamma_{ip} \\ \Gamma_{ip} = -(FV_i)^* FFV_{p-1} \\ D_p = (FV_p)^*(FV_p) = L_p L_p^* \end{cases}$$

Block ORTHODIR implementation

- Optimal solution for each RHS using all search directions computed for all RHS
- Rank revealing LL^* factorization of D_p, automatic detection of dependencies between search directions, reduction of number of search directions
- Same property as restarted ORTHODIR in term of decrease of global number of directions to be computed

- n_{block} simultaneous forward-backward substitutions at each iteration, good parallel efficiency on multi-core nodes
- Simultaneous computation of dot products, BLAS3, good parallel efficiency on multi-core nodes, global reduction for a block of scalars at once, reduced MPI overhead

- Restarted block ORTHODIR straightforward

Timings with various multi-RHS strategies

- 17x17 array ; 289 cores, 31 RHS, 50M unknowns

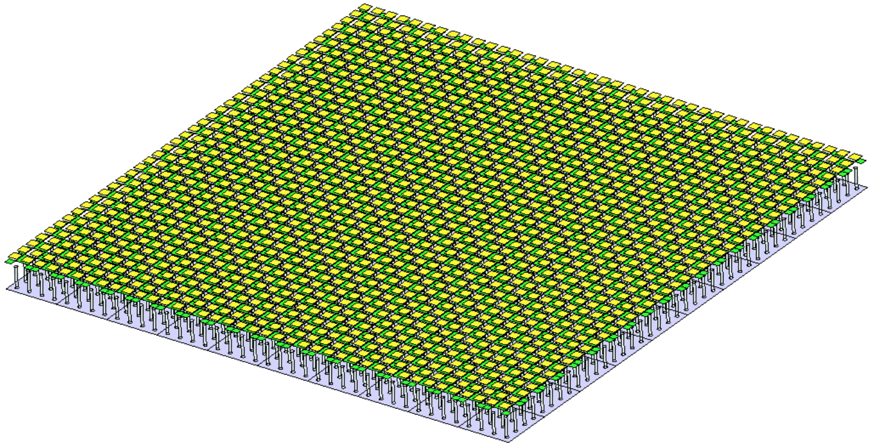

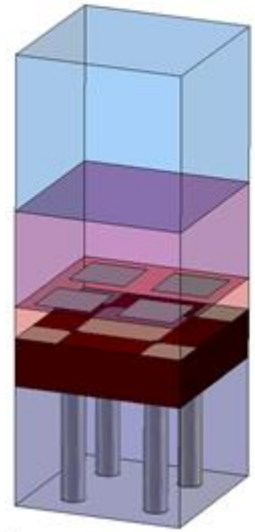

289 core 50 Million DOF	Grid	RHS	Stored directions	Elapse time(h)	Memory/core (Gb)
First RHS	17x17	1	2500	0.96	1.15
Initial RHS strategy	17x17	31	2500	3.30	1.15
Block RHS strategy	17x17	31	2500	0.8	1.4

Time / 9

Time / 37

Conclusion

Block ORTHODIR strategy is good for reducing complexity, increasing shared memory parallel efficiency of local direct solver on multi-core node and reducing message passing overhead

Make the design of good preconditioner less important